

Dynamic Testing of Xilinx Virtex-II Field Programmable Gate Array (FPGA) Input Output Blocks (IOBs)

Gary M. Swift, *Member, IEEE*, Sana Rezgui, Jeffrey George, Carl Carmichael, Matthew Napier, John Maksymowicz, Jason Moore, Austin Lesea, R. Koga, and T. F. Wrobel

Abstract— Heavy-ion irradiation and fault injection experiments were conducted to evaluate the upset sensitivity of the Xilinx Virtex-II FPGA's Input Output Block (IOB). Full Triple Module Redundancy (TMR) of the IOBs, in combination with regular configuration scrubbing, proved to be a quite effective upset mitigation method.

Index Terms— fault injection, radiation effects, reconfigurable FPGAs, single event effects.

I. INTRODUCTION

SRAM-based Field Programmable Gate Arrays (S-FPGAs) are commonly used in telecommunication, wireless, networking, video and DSP applications on the ground and are increasingly considered for space environments. In contrast to anti-fuse based FPGAs, the interconnections that define a user design are determined by configuration data stored in an array of storage latches similar to a dual-port static random-access memory, allowing dynamic reconfiguration as often as desired. Xilinx S-FPGAs were used in the Mars Exploration Rover (MER) landers to control critical pyrotechnics during the landing sequence and they are also used to control the rover wheel motors. Although S-FPGAs are highly flexible, they are sensitive to single event upsets (SEUs) from cosmic radiation. Recent work has focused on characterizing the SEU sensitivity of S-FPGAs with static and dynamic in-beam testing on the

commercial Xilinx Virtex-II family [1-3].

S-FPGAs are complex, modern devices with many configurable choices that make comprehensive radiation testing difficult and expensive. A complementary method for studying the effect of single event upsets is a simulation technique called fault injection (FI) in which errors are deliberately introduced into the FPGA configuration to see what effect they have on a targeted design.

In this paper we studied the SEU sensitivity of the input/output blocks (IOBs) used in the Xilinx Virtex-II 2V6000 FPGA using in-beam testing and fault injection with various mitigation methods to highlight the strengths and weaknesses of the two techniques. The IOBs are relatively complicated, configurable components associated with each of the several hundred general purpose pins on the Virtex-II device. In the course of this study we were also able to determine the relative effectiveness of two specific mitigation strategies which trade increased internal voting against greater use of limited input/output resources.

II. FAULT INJECTION BACKGROUND

The literature suggests four basic types of FI techniques that were considered for application here: (i) software modification at the high level description of a given design [4], (ii) the total reloading of an SRAM-based re-configurable FPGA's contents [4, 5], (iii) the partial reconfiguration of a memory segment and (iv) the use of commercial software solution such as the Jbits tool [6, 7]. Upset simulation, a primary goal of this work, is not really possible with technique (i) and technique (iv) does not allow the level of control, visibility, and flexibility that a custom approach does. However, technique (ii) has been used successfully for upset simulation targeting an earlier Xilinx Virtex family device [8] and technique (iii), while previously untried, should be similarly successful and significantly faster. The selected target device, the Virtex-II 6000, is segmented into over 2000 configuration frames. Thus, the potential increase in speed (versus the full reconfiguration FI technique) is about a factor of 1000 (because the target frame must be accompanied by a second dummy frame to "activate" the target frame).

Manuscript received July 21, 2004. Some research presented here was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration (NASA), under the NASA Electronic Parts and Packaging Program (NEPP), Code AE and with the collaboration of Xilinx Inc.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer or otherwise does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

G. M. Swift is with the Jet Propulsion Laboratory / California Institute of Technology, Pasadena, CA 91109 USA (telephone: 818-354-5059, e-mail: gary.m.swift@jpl.nasa.gov).

S. Rezgui, C. Carmichael, J. Moore and A. Lesea are with Xilinx, Inc., 2100 Logic Drive, San Jose, CA 95124 USA (telephone: 408-367-7719, email: sana.rezgui@xilinx.com).

J. George, J. Maksymowicz and R. Koga are with the Aerospace Corporation, El Segundo, CA, USA (telephone: 310-336-2089, email: jeffrey.s.george@aero.org).

M. Napier and T. F. Wrobel are with Sandia National Laboratories, Albuquerque, NM, USA (email: mpnapie@sandia.gov).

TABLE I
CHARACTERISTICS OF THE HEAVY ION BEAMS USED
AT THE TEXAS A&M UNIV. CYCLOTRON
Irradiated through flip-chip devices; silicon thinned to 200 microns.

Ion	Accelerated to MeV/amu	Energy (MeV)	LET	Range (microns)	Intervening media
Xe-129	24.8	977	56	86	In vacuum
Kr-78	40	2149	18.1	359	In air
Ar-40	24.8	623	7.5	242	In air

III. EXPERIMENTAL SET-UP

For dynamic radiation and FI experiments, a new board with an open socket for a Virtex-II 6000 device was designed by the Xilinx SEE Test Consortium¹ members and manufactured by the SEAKR Engineering. It includes two Xilinx FPGAs, their corresponding in-system programmable (ISP) configuration PROMs (XC18V04), and the required circuitry for the JTAG and SELECTMAP ports to access the FPGAs' configuration memory. The first FPGA is the Device Under Test (DUT), a radiation tolerant XQR2V6000, in an 1152-pin flip-chip fine-pitch package. It contains 144 18-Kbit block RAMs, 824 I/Os and 16,395,508 configuration bits.

The second FPGA, an XQR2V3000-FG676, has two separate monitoring functions (i) for configuration upsets and (ii) for individual channel functionality. First, it serves as the configuration monitor (also called the "service FPGA") and is used for configuration, readback and scrubbing (the transparent process of reloading the configuration bitstream so upsets are corrected) of the DUT while it's being irradiated. The number of errors scrubbed is logged continuously on a host computer. Second, it serves as a functionality monitor for the target DUT design. In this role, the service FPGA sends a data stream to each DUT channel and continuously detects and counts the number of errors in the DUT IOBs by checking the outputs against the data sent. Any mismatches detected by the service FPGA are sent to a separate host computer via a custom Visual Basic software. More details about this testing methodology are given in Ref. 3. For the testing reported here, the service FPGA feeds the channel inputs with a pseudo-random bitstream at 8 or 66 MHz and monitors the outputs for errors, counting them as they occurred.

IV. IOBS, TMR, AND THE TEST METHODOLOGY

In the Virtex II architecture, IOBs connect to a pin and are configured as either an input or an output and programmed with electrical parameters required. They consist of an "Input Block," with an optional single or double data rate (SDR or

DDR) register, an "Output Block," with an optional 3-state SDR or DDR register plus an optional 3-state high-current buffer, also with either a single or DDR register. There are several single-ended and differential I/O standard options available, including LVTTTL, LVCMOS, PCI, GTL and GTLP, HSTL, SSTL and AGP-2X standards. Additionally, a Digitally Controlled Impedance (DCI) feature provides configurable on-chip termination for each I/O element.

The target S-FPGA was configured with designs that implements multiple short "channels" of an input routed immediately to a nearby output. The simplest version of this is illustrated for one channel in Fig. 1; this was dubbed the "No TMR" design or Target Design #1. Using pins in close physical proximity minimizes the number of critical (or "used") configuration bits controlling routing resources; this maximizes the fraction of "used" bits that actually controls the IOBs' configuration.

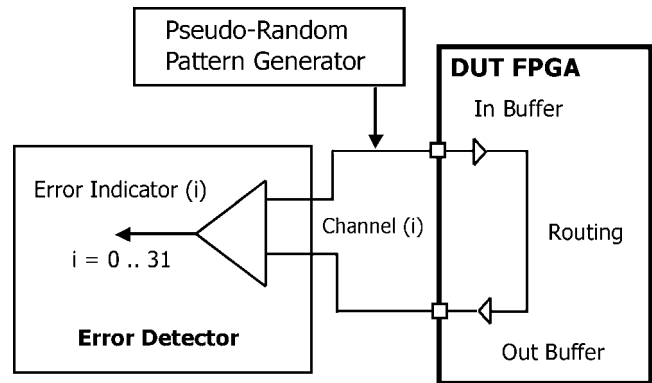


Fig. 1. Target Design 1 is the "No TMR" control case and has a single input connected with a minimum of routing to a single output and, thus, uses two external pins. Note "i=0..31" indicates that the representative channel and error detector shown are duplicated 32 times, albeit with some differences in the length of the routing constructed by the design instantiation software.

It has been shown achieving maximum S-FPGA robustness to upset requires both (a) TMR and (b) configuration scrubbing [2]. To maintain operation of a design in the presence of a single configuration upset requires a TMR design, and frequent configuration scrubblings are needed to prevent the accumulation of errors. Detailed recommendations regarding TMR implementation for the Virtex family are given in an application note [12] and apply directly to the Virtex II

¹ The Xilinx SEE Test Consortium, started by JPL and Xilinx, is a voluntary group of organizations that have a mutual interest in the evaluation of S-FPGAs for aerospace applications. The members combine resources and collaborate in order to obtain more sophisticated and efficient experimentation.

also. “Full TMR” requires the use of three pins tied together externally both for the input and for the output, as illustrated in Fig. 2; this design was dubbed Target Design #2. Three minority voters add overhead and increase the number of configuration bits involved enough so that it is not clear whether the recommended scheme is superior to a simpler output scheme: a single majority voter to a single pin. This question is definitively answered in the experimental results below by comparing with the results of “Full TMR” and Target Design #3 which is shown in Fig. 3 and was dubbed the “TMR In-only” design.

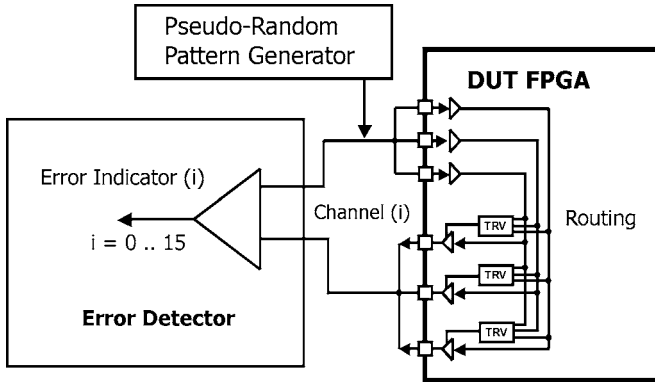


Fig. 2. Target Design 3 is the “Full TMR” case and uses three IOBs as triplicated inputs connected to three more IOBs as triplicated outputs using minimum routing and triplicated minority voters, labeled “TRV.” Thus, each channel requires six pins.

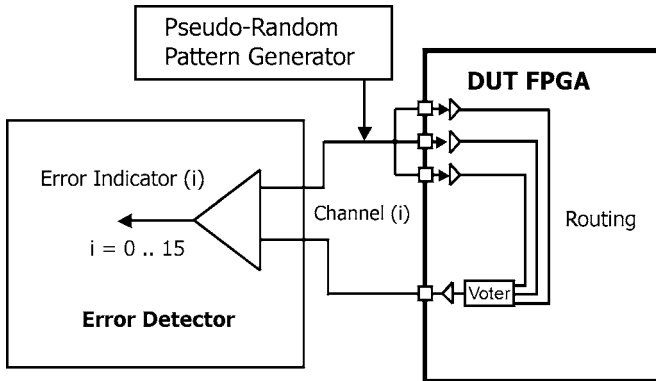


Fig. 3. Target Design 3 is the “TMR In-only” case and uses three IOBs as triplicated inputs connected to a single IOB for output using minimum routing and a single majority voter. Thus, each channel requires four pins.

V. HEAVY ION UPSET TEST RESULTS

Table I lists beam details for the accelerated heavy ions used in the upset testing. Because the target device is only available in the flip-chip configuration, the beam reaches the active region only after penetrating the silicon substrate underlying the epitaxial layer. The substrate was thinned to 200 microns and the beam energies, ranges and LETs listed in Table I are the average parameters as the ions reach the active region. The applied internal bias on the configuration bits was

2.5V in all cases and the testing was done at room temperature. The configuration scrub frequency (approximately 0.4 Hz) and the beam flux (between 23 and 350 ions/cm²/s) were set so that less than 100 errors accumulated between scrubs for most of the irradiations.

Results for irradiation of these three designs configured for the LVC MOS I/O standard at two different standard voltages for the TMR versions are given in Tables II. Fig. 4 compares the LVC MOS 3.3V results for the three target designs. As one would expect, the “No TMR” case is the worst of the three and the “Full TMR” is the best. Somewhat surprisingly, the “TMR In-only” design is only about a factor of two better than “No TMR.”




Fig. 4. Comparison of the heavy ion results for the three IOB designs for the case of LVC MOS 3.3V. Note that the “Full TMR” design has yielded zero broken channels and, therefore, only the tops of the 95% error bars can be seen here.

In order to estimate the number of configuration bits involved in each measured cross section in Table II, the parameters for static upset cross section as a function of LET for configuration bits reported in Ref. 1 were used. These yield per bit cross sections of $5.5 \times 10^{-8} \text{ cm}^2$ at LET=56 and $2.4 \times 10^{-8} \text{ cm}^2$ at LET=18.1. Assuming that these values are applicable, the effective number of configuration bits in the unmitigated channel design (row 1 of Table 2) is 91 bits at LET=56 and 35 at LET=18.1; averaging these gives about 60 critical bits for the “No TMR” case. The Weibull curves for 60 bits and 30 bits are also shown in Fig. 4; the data suggests that the IOB configuration bits measured here, while clocking data through the XQR2V6000, have a somewhat different shape from that of the static cross sections previously measured for the XQ2V1000.

Fig. 5 compares the two voltages of LCMOS “TMR In-only” results as reported in Table II. Because the configurations should be almost identical, close results were expected. The lower voltage might have some increased cross section due to lowered margins and, therefore, increased vulnerability to single-event transients and/or output “fights” due to a hit crippling one of the minority voters. Fig. 5 shows

TABLE II.
MEASURED IOB CROSS SECTIONS FOR LVCMOS I/O STANDARDS.

Voltage	TMR Design	# of channels	LET	Fluence (ions/cm ²)	Broken Channels	Cross Section Per Channel (cm ²)	# of SEFIs
3.3 V	D1- No TMR	32	56	$0.43 \cdot 10^5$	7	$5.1 (+5.4, -3.1) \cdot 10^{-6}$	0
			18.1	$35.2 \cdot 10^5$	94	$0.84 (+0.17, -0.17) \cdot 10^{-6}$	3
3.3 V	D2- Full TMR	16	56	$2.7 \cdot 10^5$	0	$0.0 (+0.86, -0.0) \cdot 10^{-6}$	1
			18.1	$35.0 \cdot 10^5$	0	$0.0 (+0.066, -0.0) \cdot 10^{-6}$	6
			7.5	$67.3 \cdot 10^5$	0	$0.0 (+0.034, -0.0) \cdot 10^{-6}$	14
3.3 V	D3- TMR In only	25	56	$5.9 \cdot 10^5$	25	$1.7 (+0.8, -0.6) \cdot 10^{-6}$	3
		16	18.1	$40.7 \cdot 10^5$	46	$0.71 (+0.24, -0.19) \cdot 10^{-6}$	5
		16	7.5	$4.0 \cdot 10^5$	11	$0.17 (+0.14, -0.09) \cdot 10^{-6}$	2
1.8 V	D2- Full TMR	16	56	$46.0 \cdot 10^5$	1	$0.15 (+0.70, -0.14) \cdot 10^{-6}$	3
			7.5	$6.1 \cdot 10^5$	0	$0.00 (+0.38, -0.00) \cdot 10^{-6}$	0
1.8 V	D3- TMR In only	25	56	$3.1 \cdot 10^5$	9	$1.2 (+1.1, -0.7) \cdot 10^{-6}$	1
		16	7.5	$3.0 \cdot 10^6$	23	$0.48 (+0.24, -0.18) \cdot 10^{-6}$	3

TABLE III.
MEASURED IOB CROSS SECTIONS FOR LVDS I/O STANDARDS

Voltage	TMR Design	# of channels	LET	Fluence (ions/cm ²)	Broken Channels	Cross Section Per Channel (cm ²)	# of SEFIs
3.3 V	D1- No TMR	16	18.1	$1.00 \cdot 10^7$	96	$0.60 (+0.12, -0.12) \cdot 10^{-6}$	20
			7.5	$9.28 \cdot 10^4$	2	$1.4 (+3.5, -1.2) \cdot 10^{-6}$	0
3.3 V	D2- Full TMR	7	18.1	$1.01 \cdot 10^7$	1	$0.042 (+0.082, -0.034) \cdot 10^{-6}$	7
3.3 V	D3- TMR In only	8	18.1	$1.03 \cdot 10^7$	75	$0.91 (+0.21, -0.21) \cdot 10^{-6}$	20
2.5 V	D1- No TMR	16	7.5	$9.6 \cdot 10^6$	24	$0.16 (+0.08, -0.06) \cdot 10^{-6}$	5

the comparison is as expected: the 1.8V results are similar to the 3.3V results - or are showing a slight (perhaps 50%) increase in susceptibility.

See Fig. 6 and Table 3 for results on the three designs implemented with low-voltage differential signaling (LVDS)

I/O standards. These runs were done with an order of magnitude higher flux of between 2100 and 4000 ions/cm²/s. This higher flux should not affect the “No TMR” results, but would tend to increase measured susceptibility of the other designs because of the strong rate dependence of upset

mitigation schemes, including TMR. Interestingly, these data suggest the “TMR-In Only” is worse than the “No TMR” case, although statistical significance is absent. They also show a significant advantage for “Full TMR,” although a single observation of three channels breaking simultaneously did occur. This is likely due to the interaction of a group of two or more configuration upsets that individually would not cause a problem if present at separate times, exactly the type of coincidence engendered by the higher flux.




Fig. 5. Comparison of the heavy ion results for the “TMR In-only” IOB design for two voltages of LVC MOS. Differences, although expected, are not significant.




Fig. 6. Comparison of the heavy ion results for the three IOB designs for the case of LVDS 3.3 V.

The last column of the tables gives the number of single-event functionality interrupts (SEFIs) observed. Observationally, SEFIs are defined here as an instance when all channels (7 to 32) stop working simultaneously. It should be noted that SEFIs are independent of design. Therefore, the cross section at a given LET can be derived from the sum of the values at that LET from the last column and the total

fluence column of Table 2, yielding, for example, at LET=56 MeV per mg/cm², eight SEFIs occurred during a total fluence of $5.8 \times 10^6 \text{ cm}^{-2}$ yielding a SEFI cross section of $1.4 (+1.3, -0.8) \times 10^{-6} \text{ cm}^2$.

VI. FAULT INJECTION EXPERIMENTS

The purpose of FI is to identify sensitive bits in the configuration bits of a particular design. A sensitive bit is defined as a configuration bit that, when upset, disrupts the normal functional operation of the design. For the designs described in Section III, that means a broken channel.

To inject and repair faults, a new module, dubbed the FI Core, was developed to reside in the service FPGA along side and independent of the configuration and functional monitors. When the FI Core injects an error, the configuration monitor is used to confirm the presence of the simulated upset and the functional monitor determines whether one (or more) channels has been broken. Although the I/O stream is being clocked at 66 MHz and FI Core clocks the configuration port at 33 MHz, the functional monitor only reports errors about 30 times a second. Thus, it represents a significant bottleneck in the performance of the current FI setup.

In order to make the problem more tractable, a software tool developed by the Xilinx Research Labs called SEUPI (Single Event Upset Probability Impact Tool) was used. This tool incorporates knowledge of the internal architecture controlled by the configuration bits and is able to identify all bits which are potentially sensitive in a particular design. The bits identified by the SEUPI tool are designated “used” configuration bits here. The SEUPI tool guarantees that the state of any single bit that it has not identified as “used” will not matter to the analyzed design. Further details of how SEUPI works can be found in Ref. 11.

SEUPI was used to identify all frames with “used” bits and the FI Core injected errors one-by-one in each of the 7872 bits in the identified frames. By observing the functional monitor, all sensitive bits can, thus, be identified.

Due to time constraints, FI experiments were performed only on the LVC MOS 3.3V versions of the three target designs. Comparisons of the SEUPI output for the LVC MOS 1.8V designs showed that the voltage change added only a few bits, so the FI results would likely be the same. SEUPI had identified 20 to 46 frames with “used” bits. The FI experiments then were run, taking about 6 to 14 hours for every bit in the identified frames to be individually toggled.

In the case of the “No-TMR” design implemented with LVC MOS 3.3V I/Os, the SEUPI tool has identified an average of 238 used bits per channel and has specified what they each configure: 55 bits control the input IOB, 63 bits the output IOB and an average of 120 bits are involved in the routing (varying from a minimum of 49 to a maximum of 371).

The following FI results were obtained for this case: only 11 bits controlling IOBs are sensitive. These eleven break down

into 3 bits (out of 55 used bits) for the input IOB and 8 bits (out of 63 used bits) for the output IOB. Only an average of 32 bits controlling routing were discovered to be sensitive, ranging from 12 to 94 bits. The results of all 32 channels are shown in Fig. 7 where SEUPI-identified “used” bits (x-axis) are plotted against FI-identified sensitive bits. Even if SEUPI had not identified the type of configuration bit, a straight-line fit shows about eleven IOB bits (looking at the y-intercept).




Fig. 7. The number of sensitive bits found by fault injection for each of the 32 channels of the “No TMR” design plotted as a function of the number of “used” bits for each channel reported by the SEUPI design analysis tool. The points should fall on the line shown if there are 11 (the intercept) sensitive bits in the IOB pair and if 25% of the routing bits are sensitive.

The “Full-TMR” design implements 16 IOB channels. SEUPI identified that 165 (3 x 55) bits are used to define the triplicated input IOBs, 189 (3 x 63) bits for the triplicated output IOBs and an average of 568 bits (ranging from 482 to 640) are used for routing. Fault injection results on the “Full-TMR” design prove that no single mis-configured bit can defeat the functional operation of a fully mitigated IOB channel.

The “TMR-In Only” design also implements 16 IOB channels. In this case, the SEUPI tool identified 228 used bits (165 allocated for the triplicated input IOBs and 63 bits for the single output IOB) in addition to an average of 202 routing bits (ranging from 178 to 230 bits). FI results show that no bits controlling input IOB are sensitive. On the output IOB side, the previously identified 8 bits sensitive in the “No-TMR” design, were again found to be classified as sensitive. This is as expected because only the input IOB was mitigated in this design. In addition, an average of 11 routing bits for each channel has been found to be disruptive of its functionality. This led to an average of 19 sensitive bits on every IOB channel, which is approximately the half of the counted sensitive bits for a “No-TMR” IOB channel (43 bits). This factor of two correlates well with the radiation results of the previous Section.

VII. DISCUSSION: COMPARING FAULT INJECTION AND HEAVY ION RESULTS

This paper presents results for two complementary

methodologies for SEU testing: in-beam upset results and fault injection experiments. The DUT is the XQR2V6000 from the Virtex-II S-FPGA family, and the example results focus on the input/output circuitry. The former (radiation testing) was tailored to directly measure the IOBs’ upset cross-sections. The latter (fault injection) was used to identify the precise result of particular bit errors. That the two techniques give similar results can be seen in Table IV below, although it appears that the beam results are about 50% higher. This discrepancy may be due to the effects of single-event transients in the data stream; the beam can cause transients and fault injection cannot. Also, either the occurrence of coincident errors from multiple ion strikes within a scrub window or multiple upsets from a single ion strike may explain the difference. These hypotheses will be sorted out using a new capability being developed that uses a communication channel between the configuration monitor and the functional monitor. When an interesting event, such as a channel’s functionality is broken, the functional monitor will signal the configuration monitor to forego scrubbing momentarily and capture the current configuration for later analysis. If SEUPI and/or FI can confirm that none of the upset bits in the configuration alone can account for the broken channel, then FI will be used to determine if any pairs of bits are responsible. If not, then by elimination, the transient explanation is the correct one. It is interesting to note that the new in-beam capture capability is really only useful in conjunction with fault injection.

VIII. CONCLUSIONS

Two specific conclusions can be drawn here. First, the recommended TMR using three input pins and three output pins is very effective at mitigating functional problems caused by a configuration upset (in conjunction with configuration scrubbing). Second, “partial” TMR on I/Os, such as the “TMR In-only case,” adds only a small increment of upset mitigation.

TABLE IV.
COMPARISON OF THE COUNTS OF SENSITIVE BITS FROM FAULT INJECTION
AND BEAM UPSETS

3.3V LVCMOS Design	FI	Beam
No TMR	43	60
Full TMR	0	0
TMR-In Only	19	30

General conclusions may be reached inductively from the IOB beam and FI investigations. First, the fault injection technique appears most useful as an adjunct to upset testing where additional or more precise information needs to be developed. Thus, the two techniques are complementary. While FI can reduce the beam time needed for upset characterization, the results of FI need some irradiation data for calibration and also to validate that upsets that cannot be simulated by faults do not dominate in-beam results. Second,

the use of partial reconfiguration allows faults to be injected relatively quickly; this results in the speed of the functional testing limiting the speed of fault injection. Finally, fault injection cannot fully replace heavy ion irradiation, not only because threshold LET must be found with the beam, but also because the beam can induce transients and FI cannot; transients do appear to be important at the speeds of modern devices.

REFERENCES

- [1] C. Yui, G. Swift and C. Carmichael, "Single Event Upset Susceptibility Testing of the Xilinx Virtex II FPGA," presented at 2002 Int'l Conf. on Military and Aerospace Applications of Programmable Logic Devices, Sept. 2002.
- [2] C. Yui, G. Swift, C. Carmichael, R. Koga and J. George, "SEU Mitigation Testing of Xilinx Virtex II FPGAs," *2003 Radiation Effects Data Workshop Record*, pp. 92 - 97.
- [3] R. Koga, J. George, G. Swift, C. Yui, C. Carmichael, T. Langley, P. Murray, K. Lanes and M. Napier, "Comparison of Xilinx Virtex-II FPGAs SEE sensitivities to Protons and Heavy Ions," *Proc. 7th European Conf. on Radiation and Its Effects on Components and Systems*, Sept. 2003.
- [4] T. Delong, B. Johnson and J. Profeta, "A Fault Injection Technique for VHDL Behavioral-Level Models," *IEEE Design Test Computing*, Vol. 13, pp. 24-33, Winter 1996.
- [5] F. Lima, C. Carmichael, J. Fabula, R. Padovani and R. Reis, "A Fault Injection Analysis of Virtex FPGA TMR Design Methodology," *Proc. 6th European Conf. on Radiation and Its Effects on Components and Systems*, pp. 275 - 282, Sept. 2001.
- [6] L. Antoni, R. Leveugle and B. Feher, "Using Run-Time Reconfiguration for Fault Injection Applications," *IEEE Trans. on Instrumentation and Measurement*, Vol. 52, No. 5, Oct. 2003.
- [7] M. Caffrey, P. Graham, E. Johnson and M. Wirthlin, "Single-Event Upsets in SRAM FPGAs," presented at 2003 Int'l Conf. on Military and Aerospace Applications of Programmable Logic Devices, Sept. 2002.
- [8] M. Alderighi, F. Casini, S. D'angelo, M. Mancini, A. Marmo, S. Pastore, G. Sechi, "A tool for injecting SEU-like faults into the configuration control mechanism of Xilinx Virtex FPGAs," *Proc. 18th IEEE Int'l Symp. on Defect and Fault Tolerance in VLSI Systems*, Nov. 2003.
- [9] G. Swift, C. Yui and C. Carmichael, "Mitigating Upsets in SRAM-based FPGAs from the Xilinx Virtex 2 Family," presented at 2003 Int'l Conf. on Military and Aerospace Applications of Programmable Logic Devices, Sept. 2003.
- [10] P. Graham, M. Caffrey, M. Wirthlin, E. Johnson and N. Rollins, "SEU Mitigation for Half-Latches in Xilinx Virtex FPGAs," *IEEE Trans. Nucl. Sci.*, v. 50, no. 6, pp. 2139 - 2146, Dec. 2003 Pages:.
- [11] P. Sundararajan, C. Patterson, C. Carmichael, S. McMillan and B. Blodget, "Estimation of Single Event Upset Probability Impact of FPGA Designs," presented at 2003 Int'l Conf. on Military and Aerospace Applications of Programmable Logic Devices, Sept. 2003.
- [12] C. Carmichael, "Triple Module Redundancy Design Techniques for Virtex FPGAs," Xilinx Application Note XAPP197, Nov. 2001, available at www.xilinx.com/bvdocs/appnotes/xapp197.pdf.